Towards exa-scale ocean simulation

Oriol Tintó Prims

Advisors:

- Anna Cortés
- Mario C. Acosta
- Francisco J. Doblas-Reyes

26/04/2019

Universitat Autònoma de Barcelona

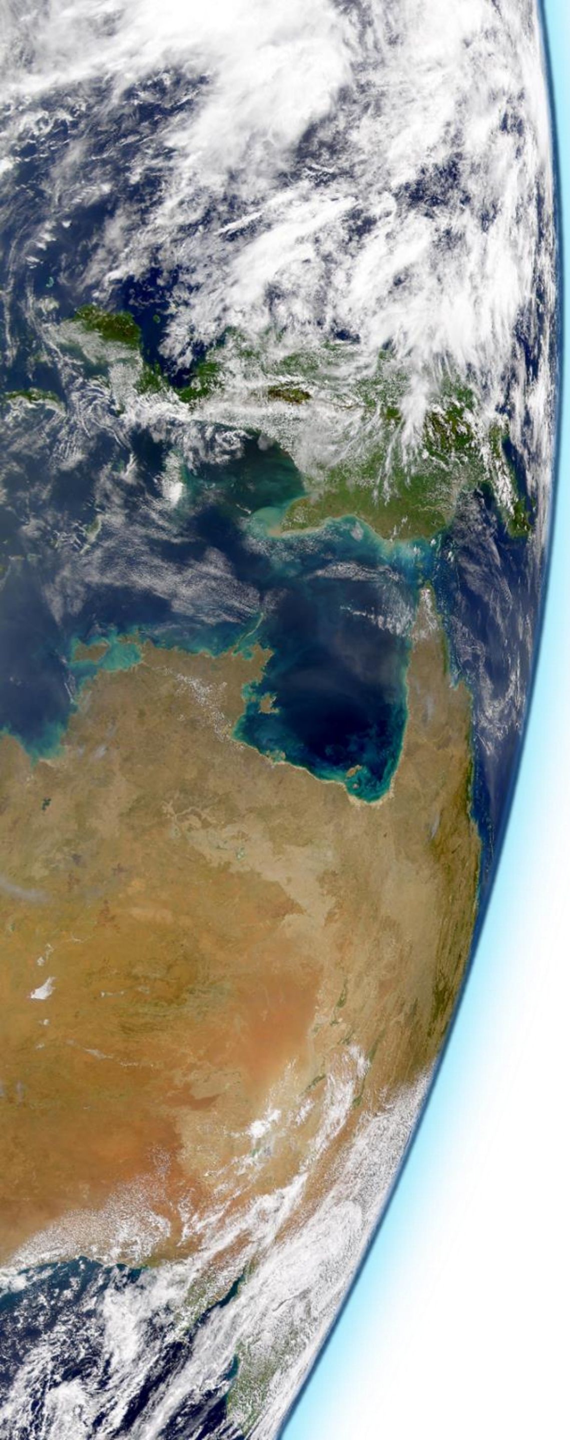

Outline

- Introduction
- Challenges
- Work done
- Future Plans

Motivation

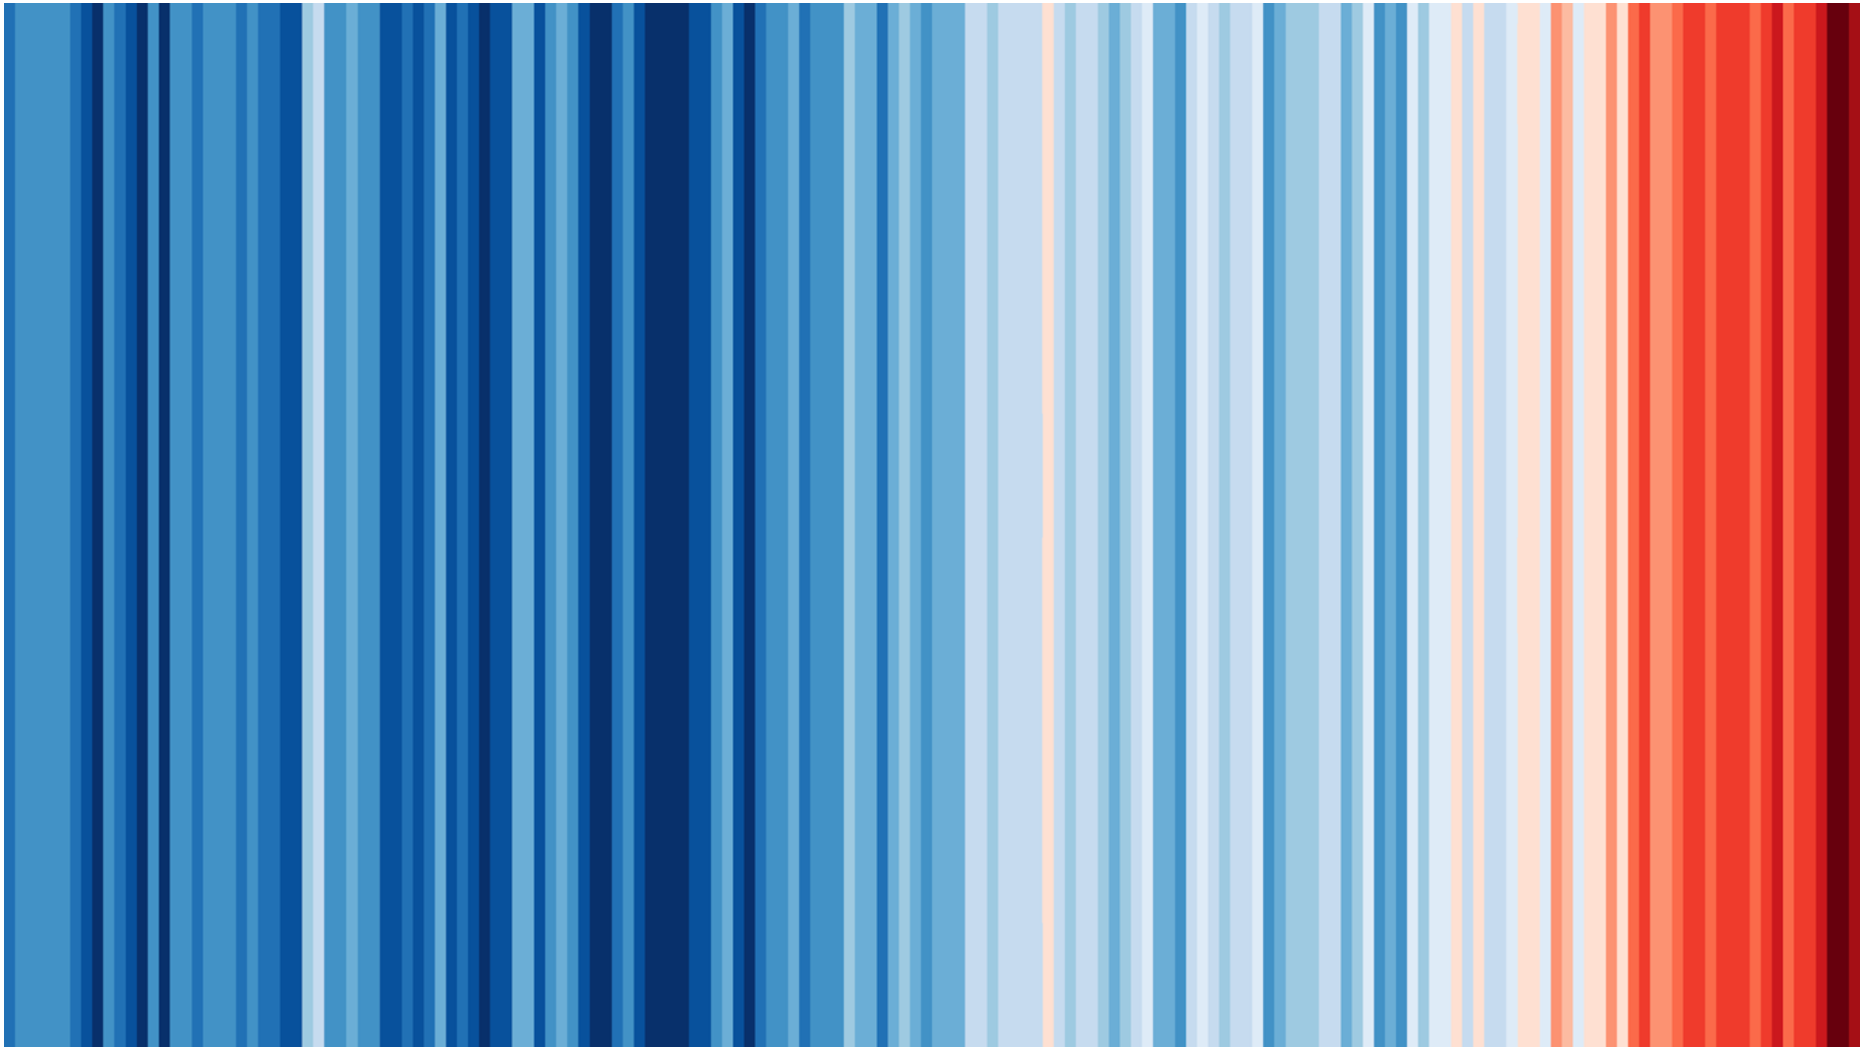

Source: **@ed_hawkins**

Motivation


IOP Publishing Environ. Res. Lett. 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

CrossMark

OPEN ACCESS

RECEIVED 28 April 2015

REVISED 27 November 2015

ACCEPTED FOR PUBLICATION 1 March 2016

PUBLISHED 13 April 2016

REPLY

Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], N
Ed W Maibach[9], J
Dana Nuccitelli[7], F

Climate ch

There will alw
as complex as
strong evidenc
occurring[1]. Th
of rising surfac
temperatures
average global
to many physi
most of the w
to human acti
led to changes

The existence
vital to life on
temperatures
than they are
atmospheric c
including carb
and nitrous ox
Carbon dioxid
1750 to over 3
levels that can
years). Increas
temperatures
approximately
century. The I
(IPCC) project
temperatures
centigrade de
levels, by 210

Reduce th

The scientific
sufficiently cle
is vital that all
can take now,
reduction in n

Action taken now to reduce significantly the build-up of greenhouse gases in the atmosphere will lessen the magnitude and rate of climate change. As the United Nations Framework Convention on Climate Change (UNFCCC) recognises, a lack of full scientific certainty about some aspects of climate change is not a reason for delaying an immediate response that will, at a reason cost, prevent dangerous anthropogenic interference w the climate system.

- Climate is changing.

that

strophic
ankind.

Motivation

- Projections
- Impact analysis
- Adaptation to climate change.

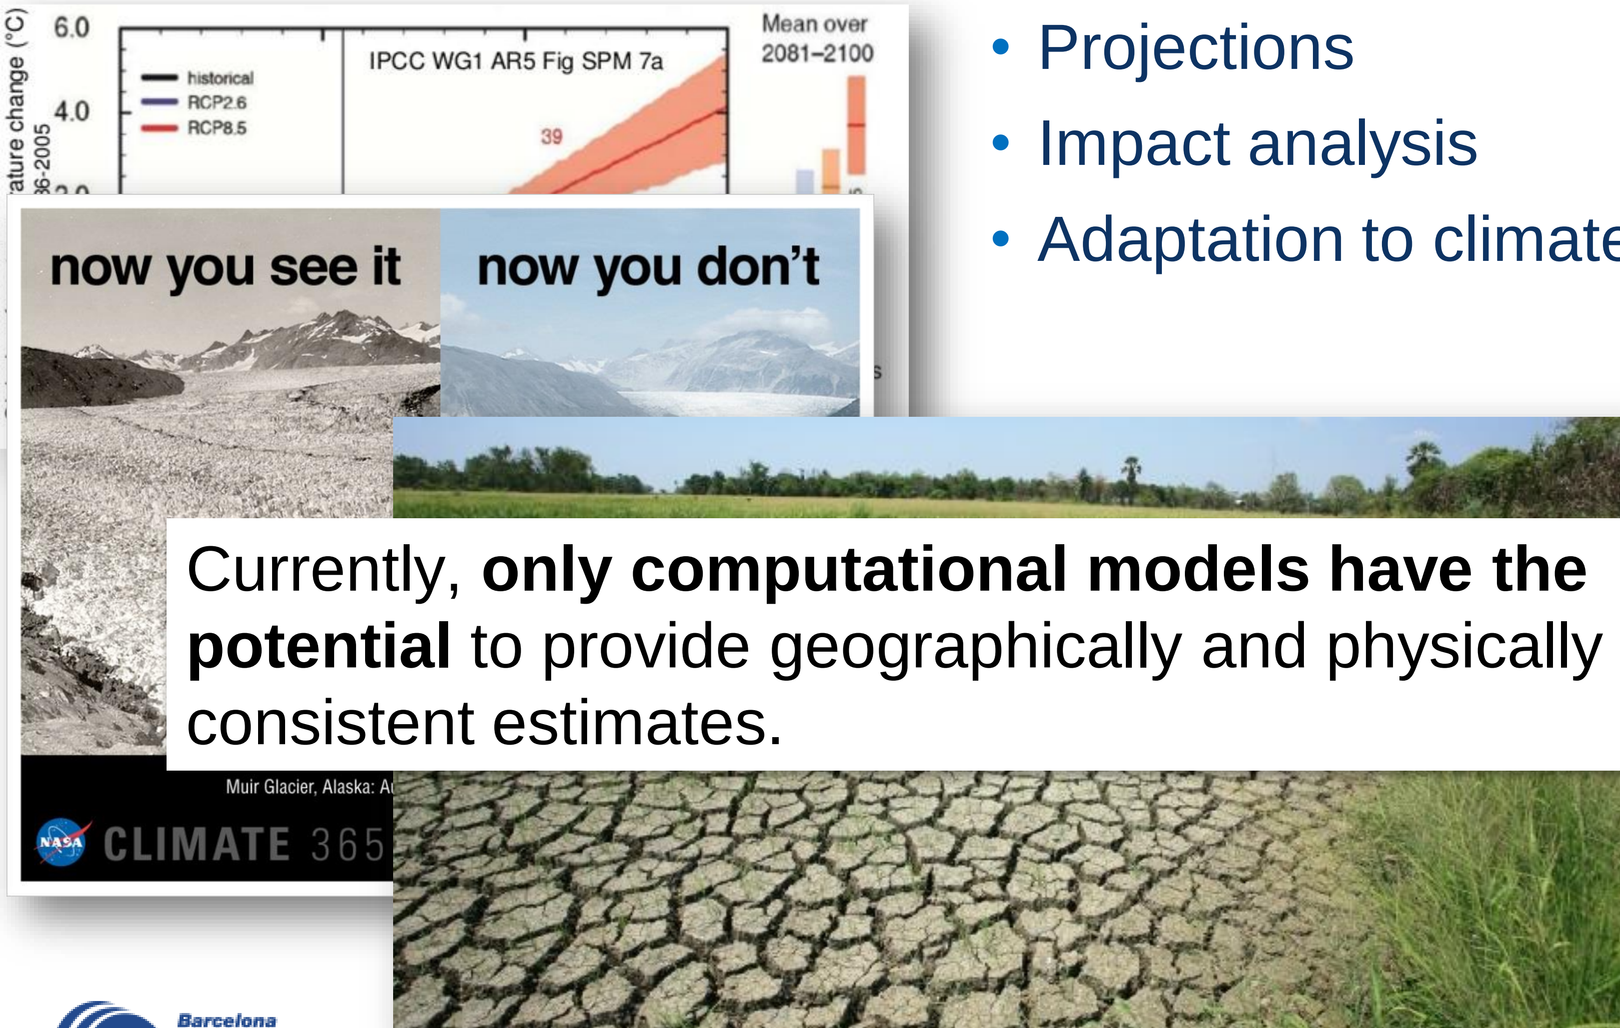

Currently, **only computational models have the potential** to provide geographically and physically consistent estimates.

Motivation

- The same models are used for both **weather and climate**.
- Weather models have a **huge value for society**.
- Improving the computational performance of these models pays back in several ways:
  - Saving resources.
  - Allowing new experiments:
    - Improving climate knowledge
    - Improving weather predictions

Nucleus for European Modeling of the Ocean (NEMO) is a **state-of-the-art** global **ocean model**

It is used in oceanographic research, operational oceanography, seasonal forecast and climate studies

Includes several **sub-models**. Many of them can work in standalone version , many others need to be coupled

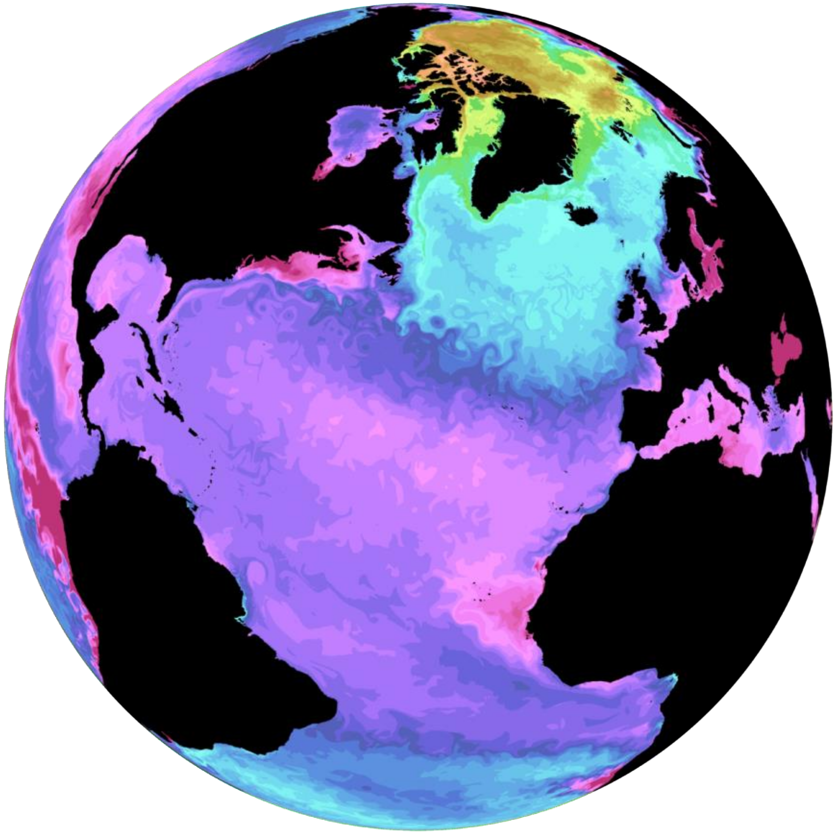

Sea Surface Temperatur e

Current Velocity

Chlorophyll

A climate modeling Timeline

Inclusion of new components

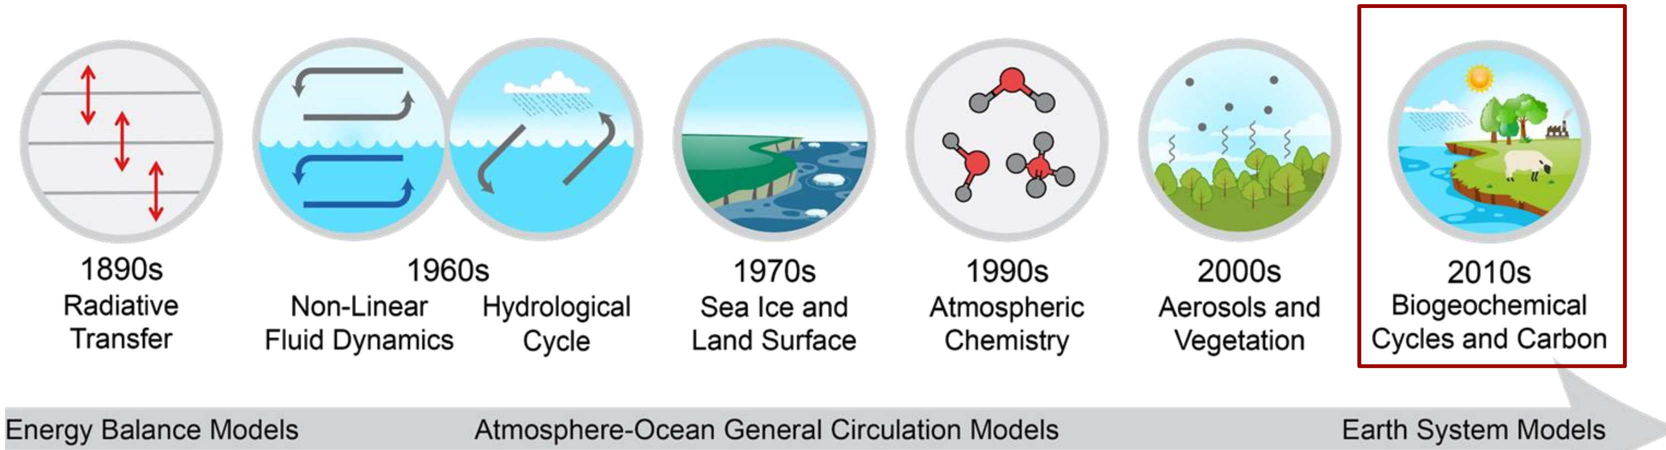

From the 4th National Climate Assessment (US), Volume I

- Allowed the representation of new climate and biogeochemical processes
- Improved the ESMs ability to represent the real world
- Provides a new framework to investigate the interactions between the different components

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos

UAB Universitat Autònoma de Barcelona

A climate modeling Timeline

Increase in spatial resolution: Ocean

Standard Resolution (1°)

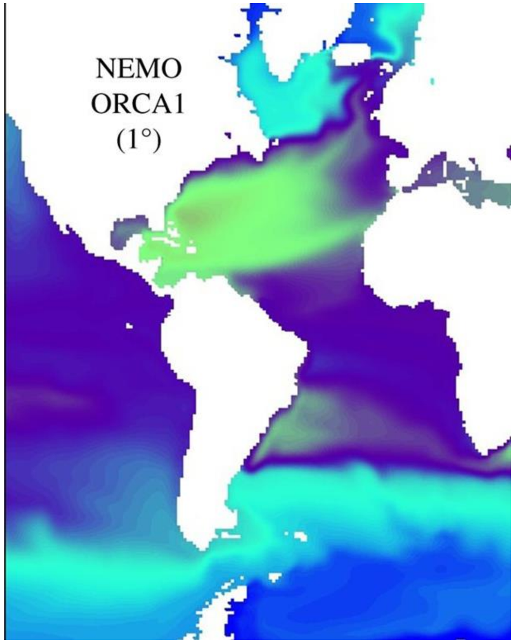

High Resolution (0.25°) Eddy-permitting

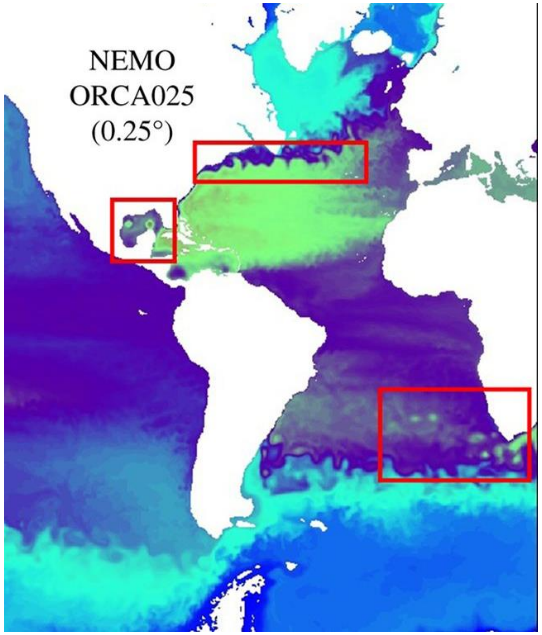

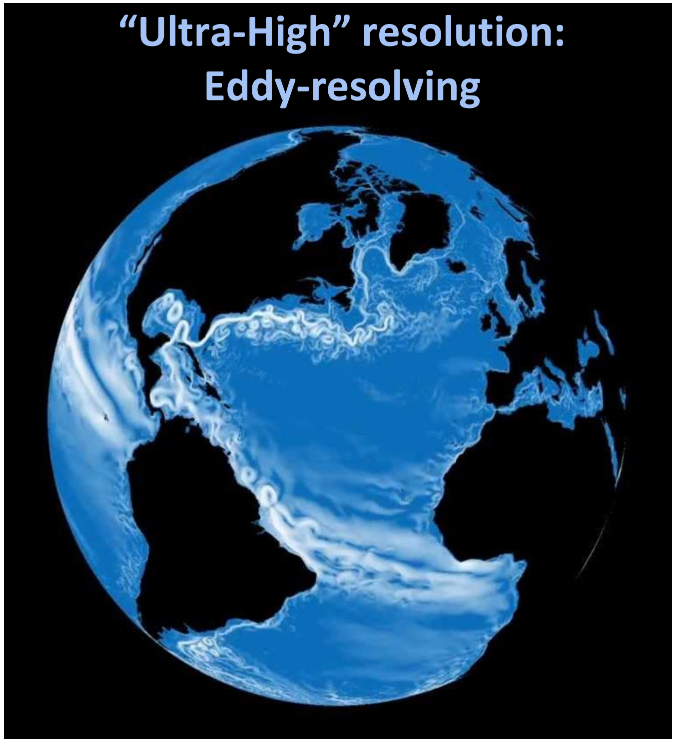

The improvements in ocean resolution translate in a better representation of eddies and ocean currents, which are key to describe realistically decadal variability in the ocean

Challenges

TOP500 Ranking of most powerful supercomputers

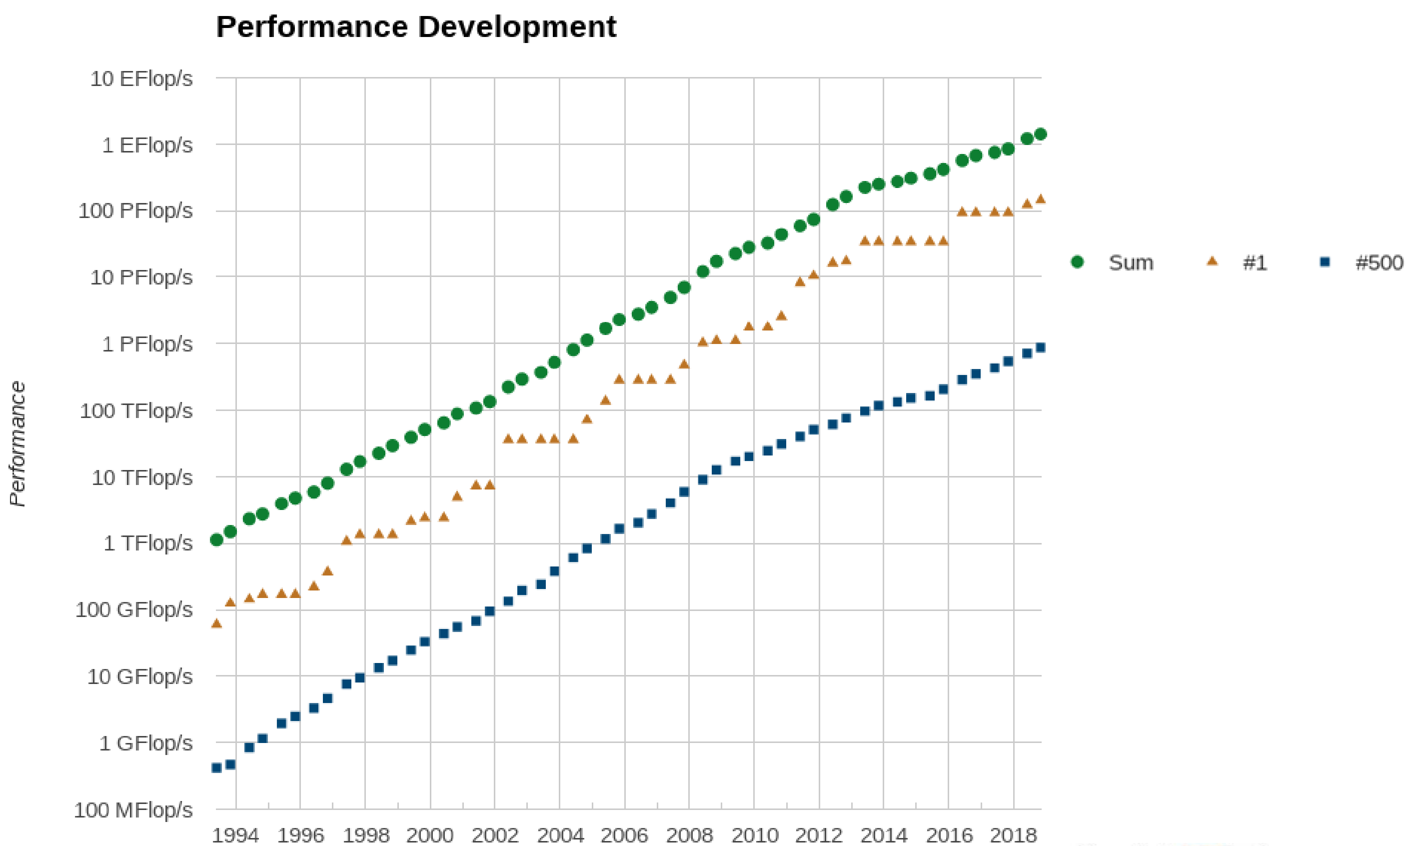

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

Challenges

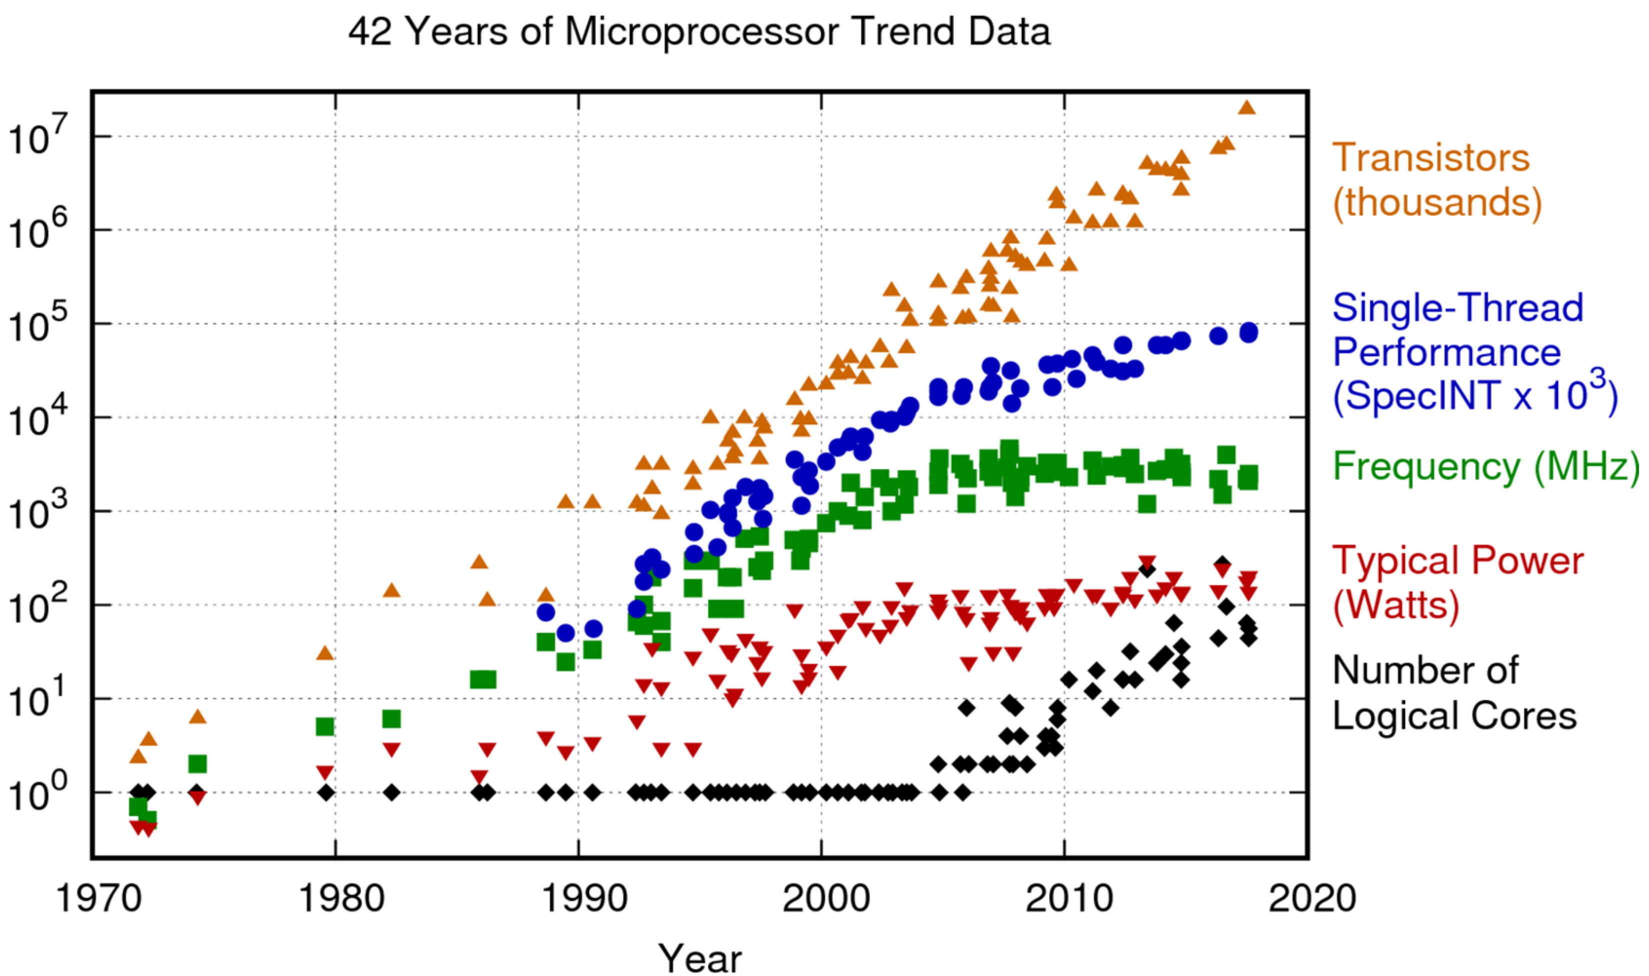

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Challenges

- How to exploit exascale systems:
  - Synchronization-reducing algorithms.
  - Communication reducing algorithms.
  - Mixed Precision methods.
  - Autotuning.
  - Fault resilient algorithms.
  - Reproducibility.

Challenges

- Start from scratch to use better the resources?

  Development cost of a climate model is estimated to be between 500-1000 person year.

  Most of models are community models and communities are reluctant to deep changes.

- Adapt the existing state-of-the-art models seems wiser choice!

Stress Test

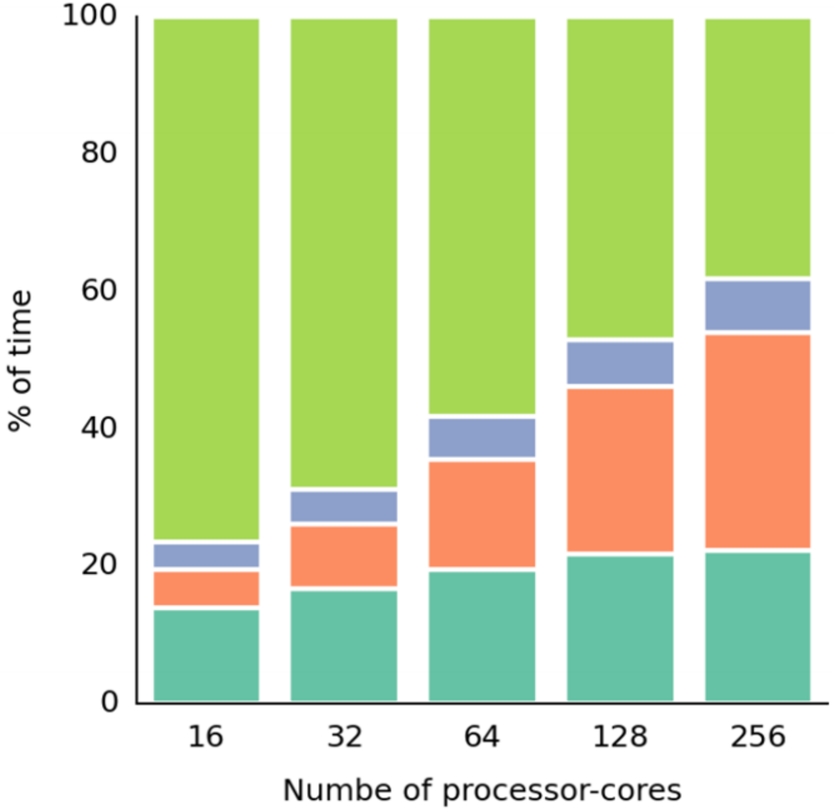

100
80
60
40
20
0
% of time
16
32
64
128
256
Numbe of processor-cores

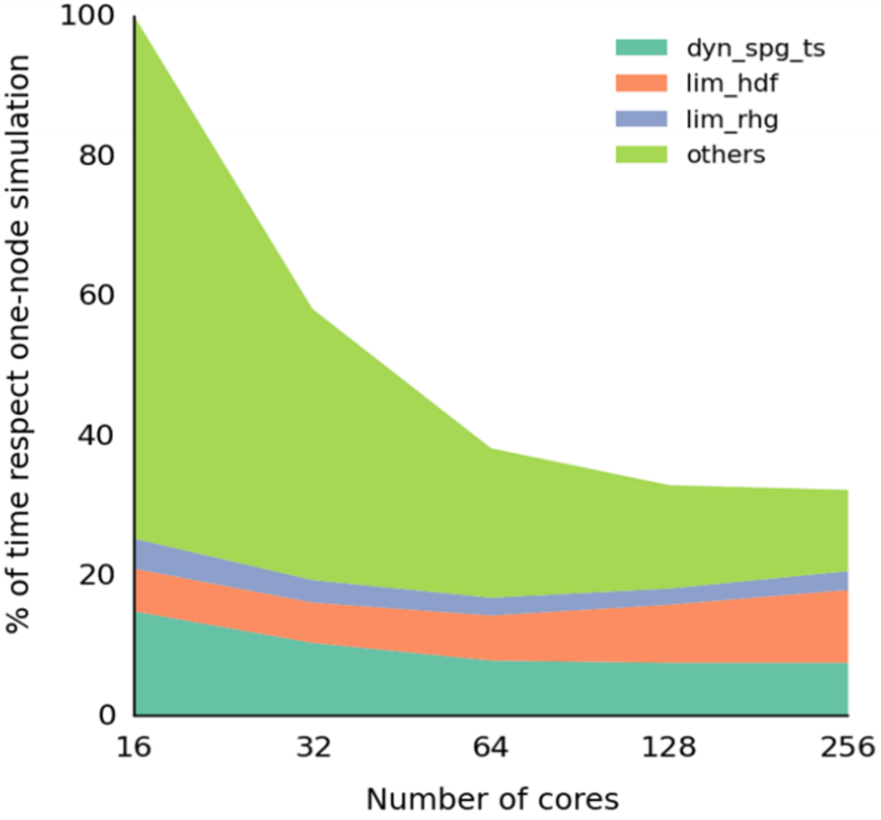

100
80
60
40
20
0
% of time respect one-node simulation
dyn_spg_ts
lim_hdf
lim_rhg
others
16
32
64
128
256
Number of cores

Impact of Communication

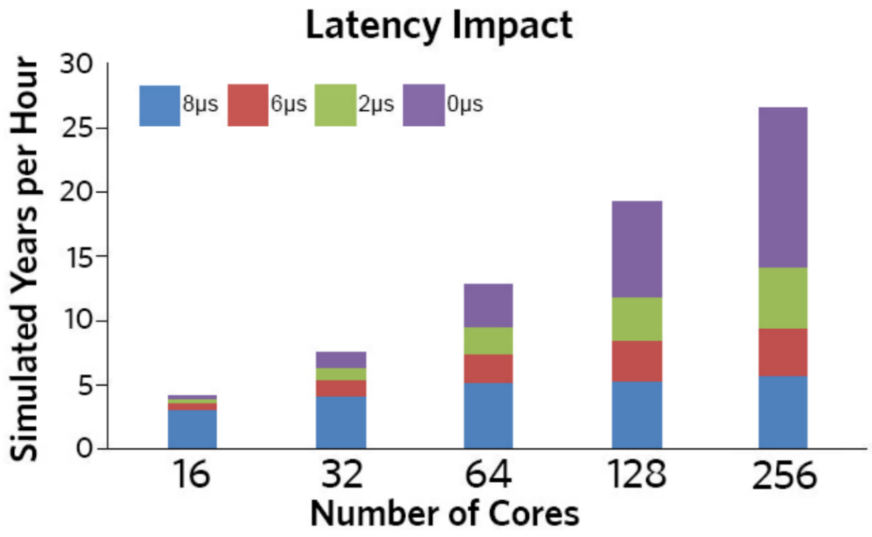

Optimizations

- Reduction of communication:
  - Message aggregation
  - Removing unnecessary collective communication

Optimizations

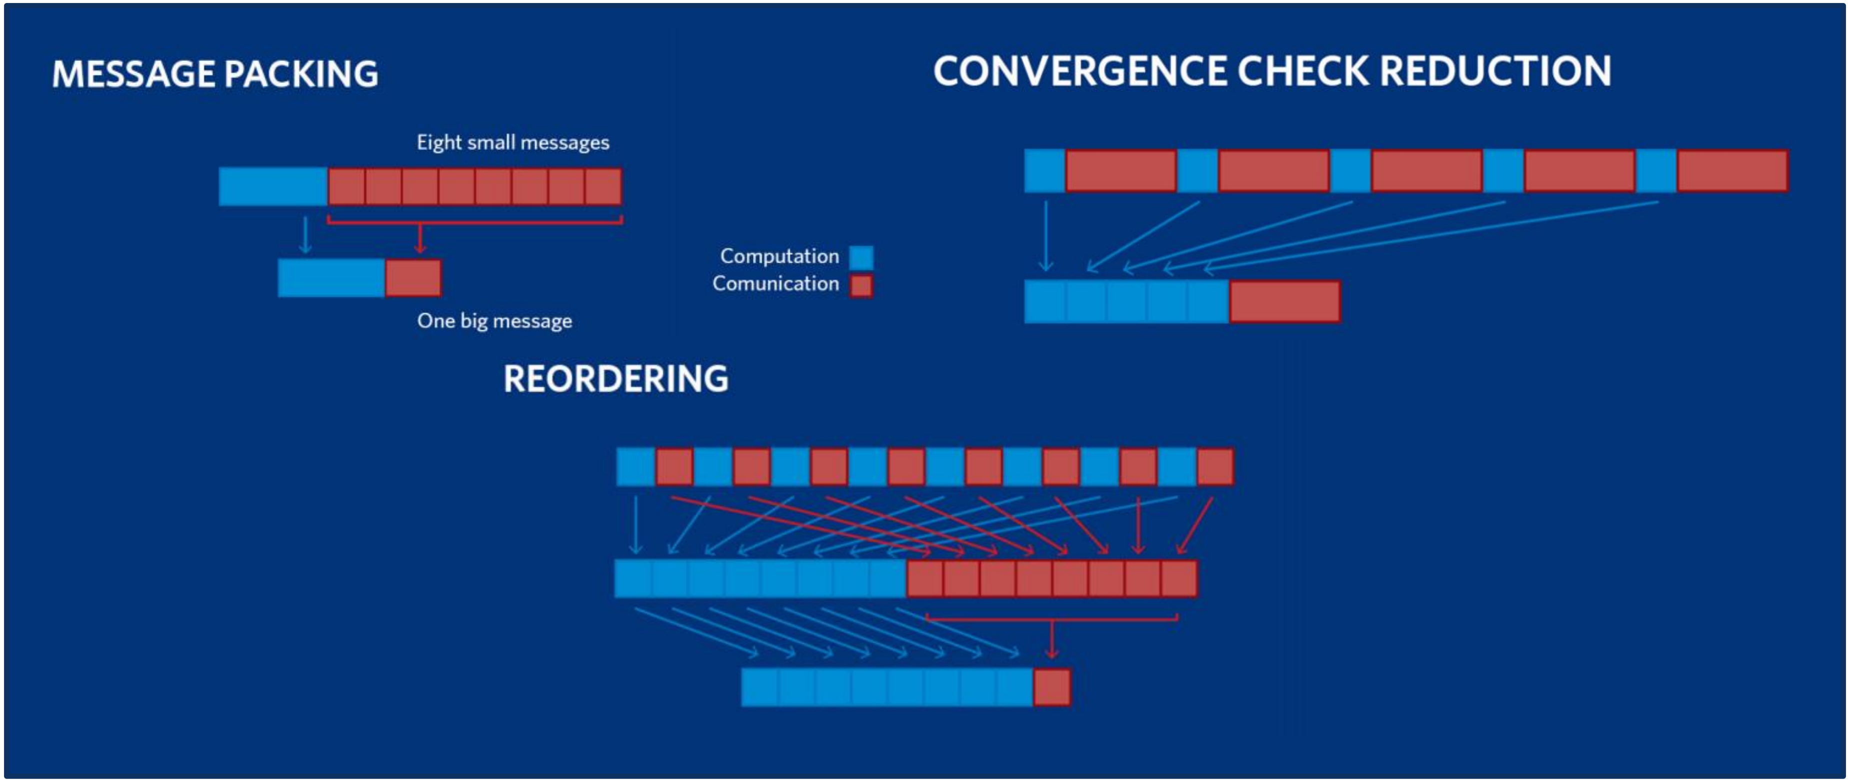

Optimization Impact

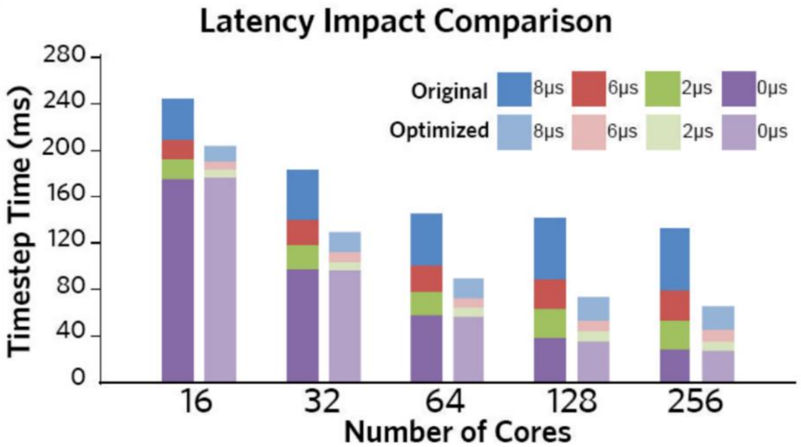

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

Optimization Impact on ORCA025 simulations

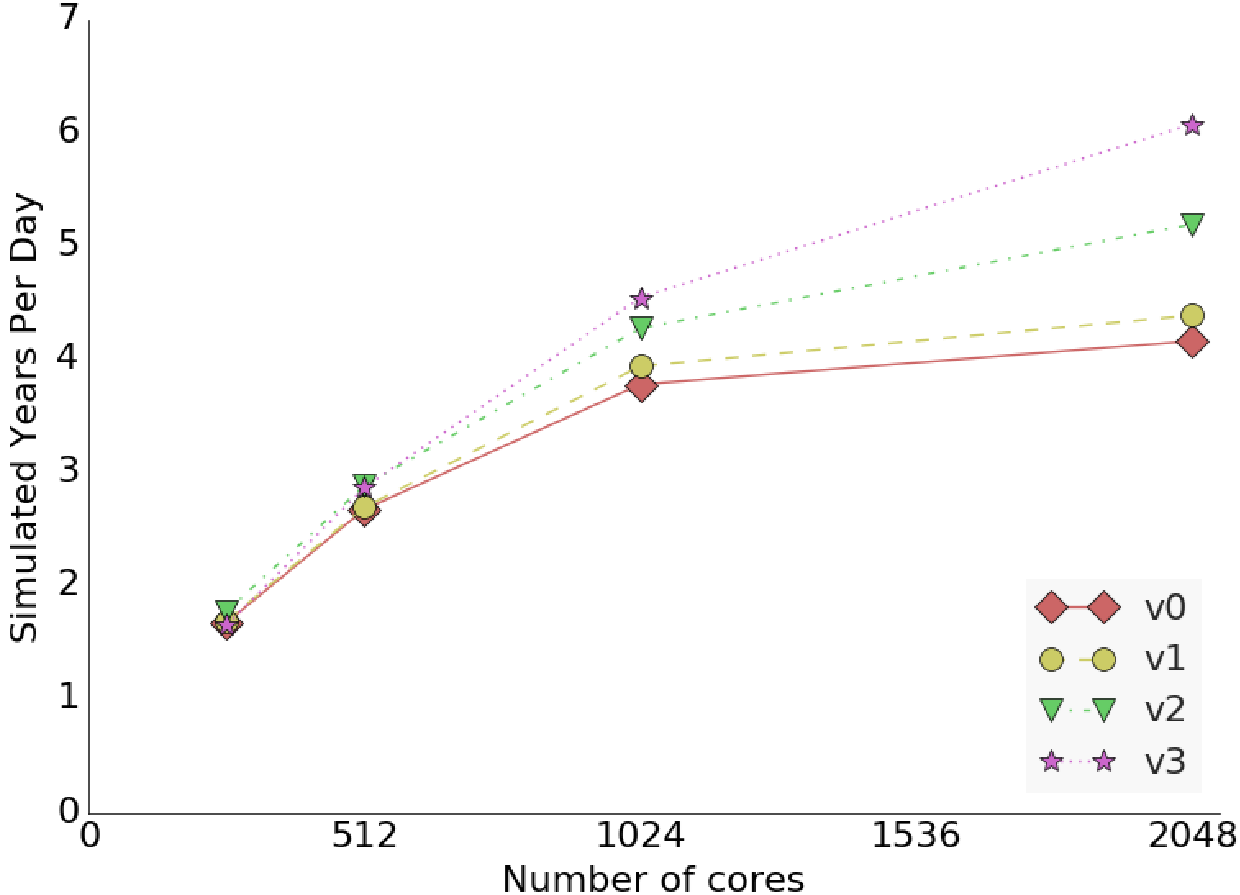

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
UAB Universitat Autònoma de Barcelona

Impact of our work

- Our optimizations have been included and used from NEMO 3.6 . Hundreds of users around the world take profit of it and millions of computer hours are saved.

- Publication:
  - Finding, analysing and solving MPI communication bottlenecks in Earth System models - Journal of computational Science (https://doi.org/10.1016/j.jocs.2018.04.015)

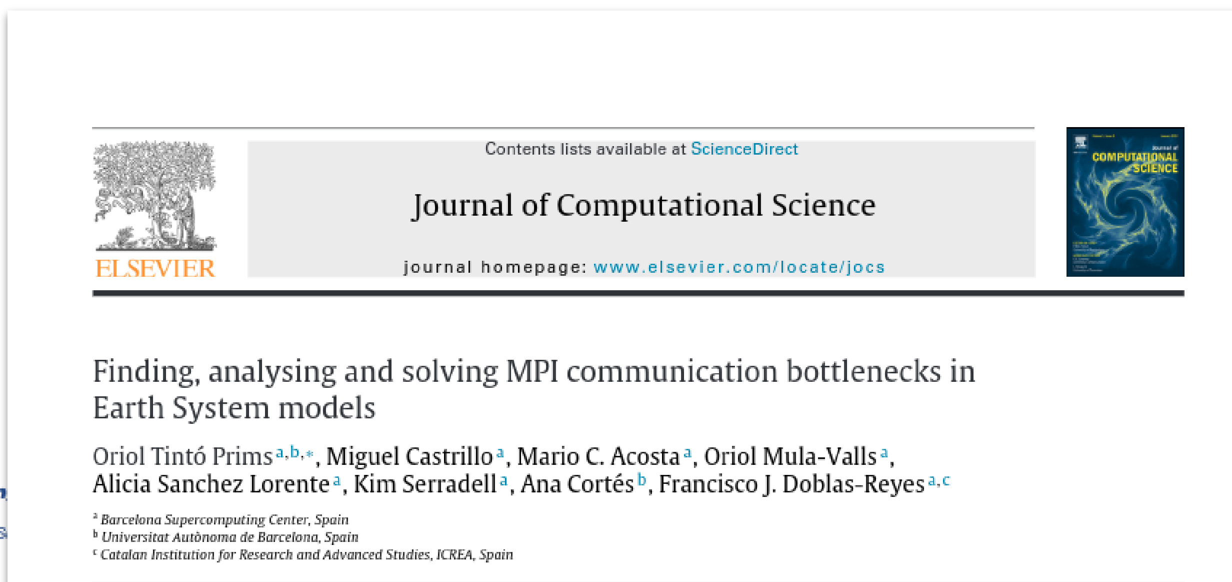

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

ELSEVIER

Finding, analysing and solving MPI communication bottlenecks in Earth System models

Oriol Tintó Prims[a,b,*], Miguel Castrillo[a], Mario C. Acosta[a], Oriol Mula-Valls[a], Alicia Sanchez Lorente[a], Kim Serradell[a], Ana Cortés[b], Francisco J. Doblas-Reyes[a,c]

[a] Barcelona Supercomputing Center, Spain
[b] Universitat Autònoma de Barcelona, Spain
[c] Catalan Institution for Research and Advanced Studies, ICREA, Spain

Barcelona Supercomputing Center Centro Nacional de S

itat Autònoma Barcelona

ELPiN

- Removing the land-only processes in the smart way.

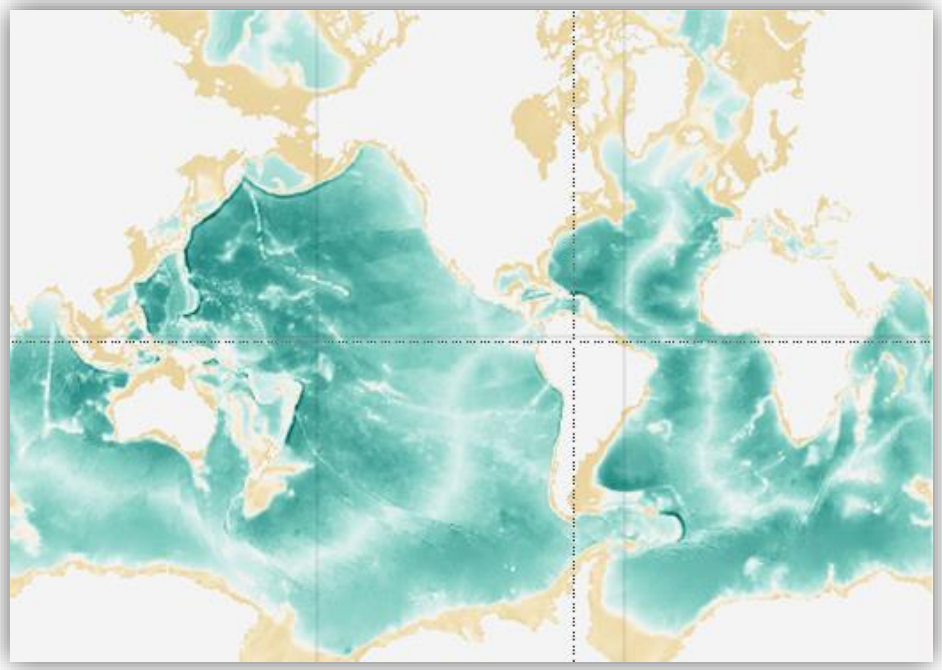

ELPiN

- Removing the land-only processes in the smart way.

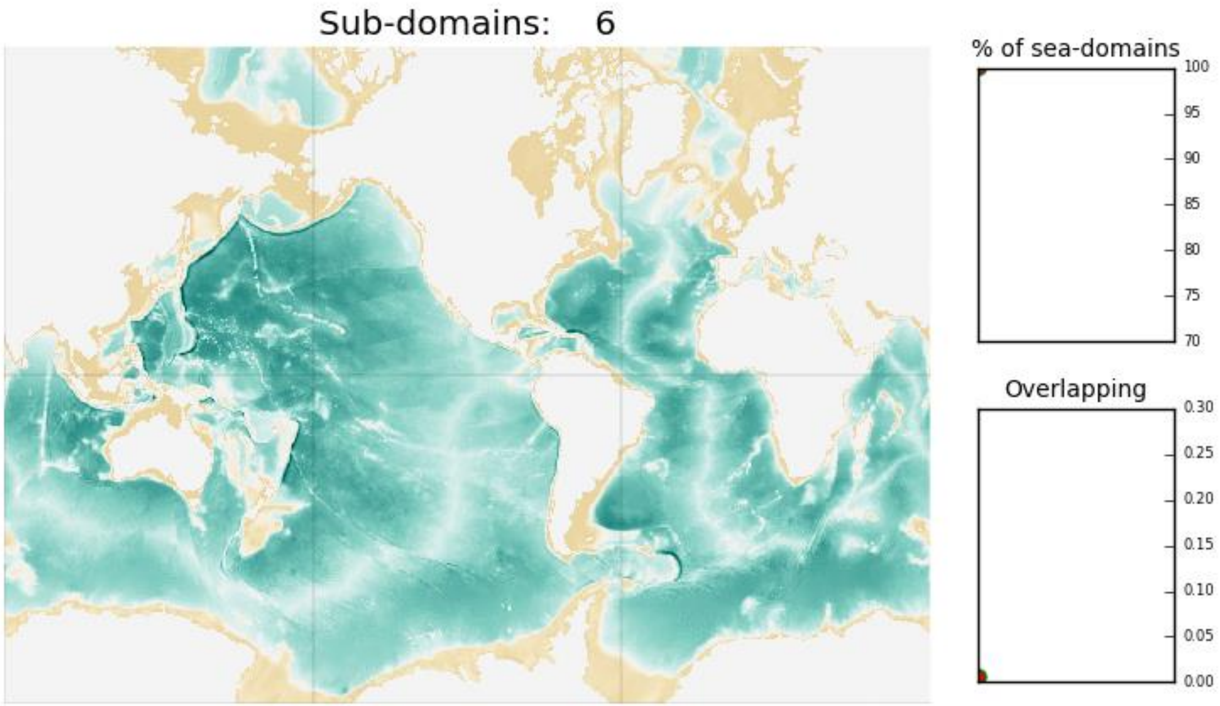

ELPiN

- Impact on ORCA025 simulations

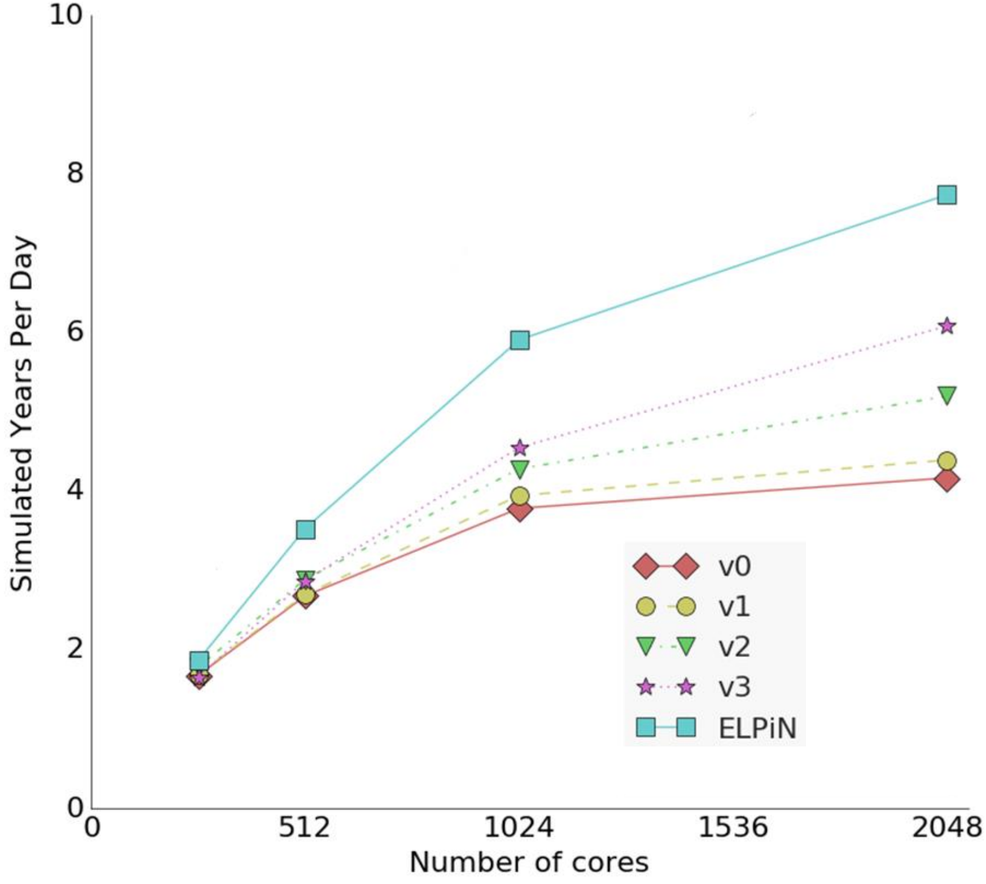

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

caos
UAB Universitat Autònoma de Barcelona

Impact of our work

- This tool has been implemented in the EC-Earth **production** workflow. CMIP-6 simulations with EC-Earth will be using it and therefore saving millions of computing hours.
- Presented and published in proceedings of the **ICCS 2017** @ Zurich

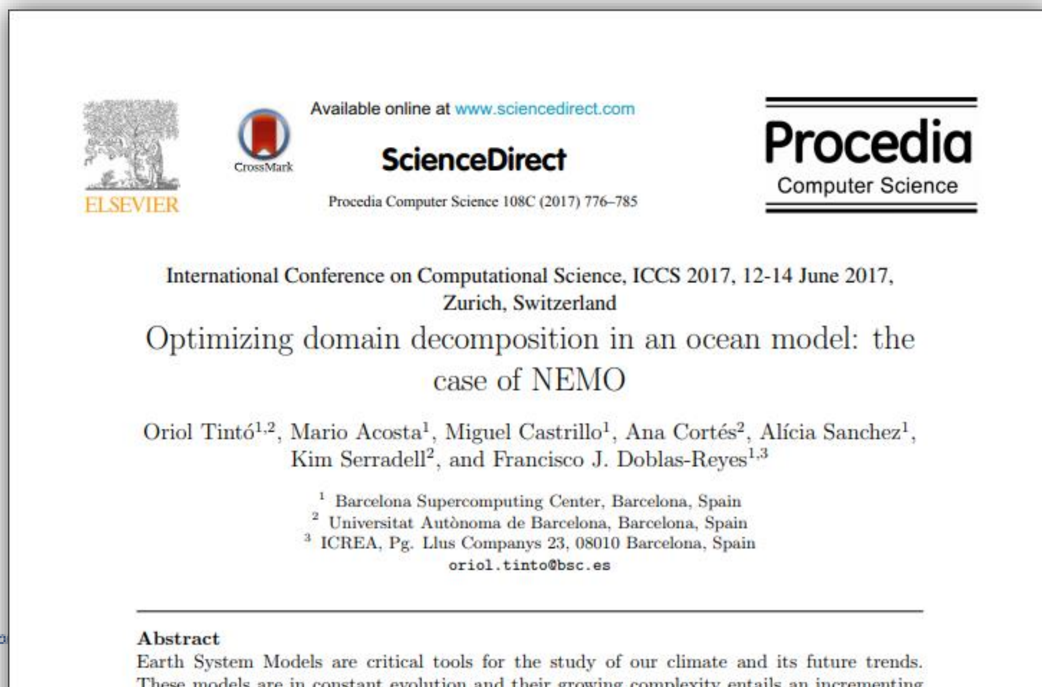

Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 776–785

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

Optimizing domain decomposition in an ocean model: the case of NEMO

Oriol Tintó[1,2], Mario Acosta[1], Miguel Castrillo[1], Ana Cortés[2], Alícia Sanchez[1], Kim Serradell[2], and Francisco J. Doblas-Reyes[1,3]

[1] Barcelona Supercomputing Center, Barcelona, Spain
[2] Universitat Autònoma de Barcelona, Barcelona, Spain
[3] ICREA, Pg. Llus Companys 23, 08010 Barcelona, Spain
oriol.tinto@bsc.es

Abstract
Earth System Models are critical tools for the study of our climate and its future trends.

How to keep improving the model?

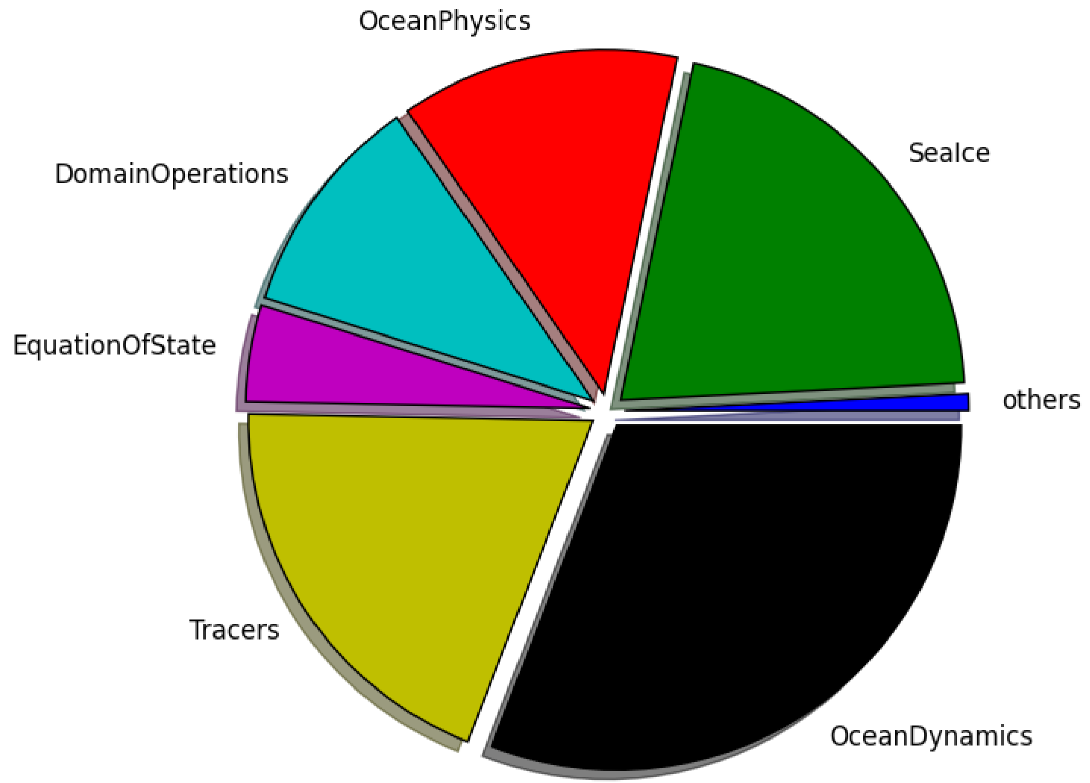

Exploring the use of mixed Precision

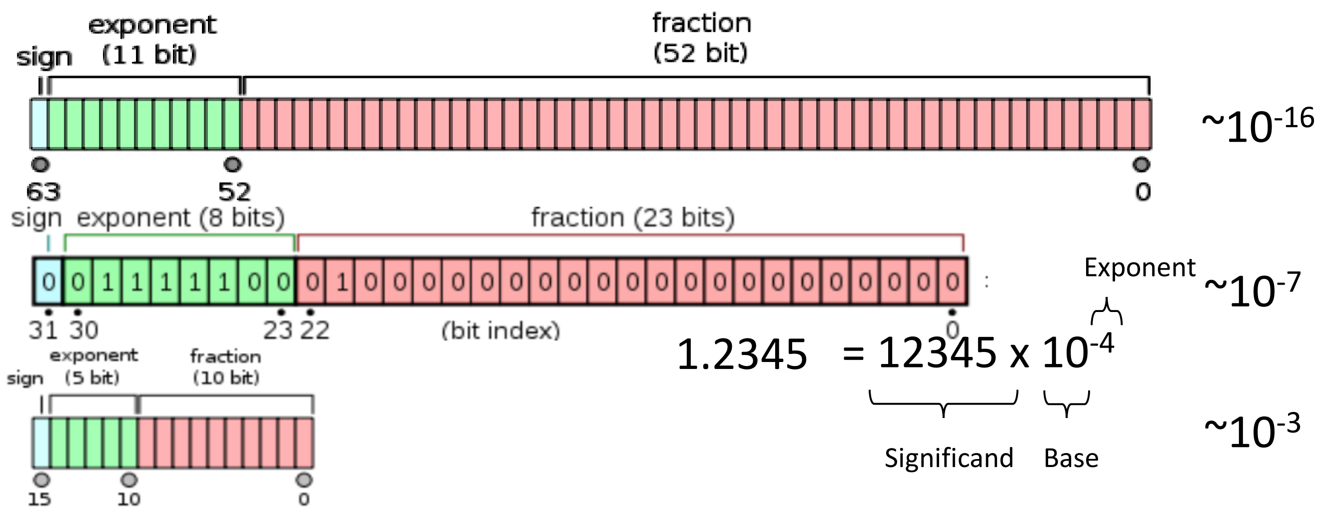

- Which precision is needed for the data that we want to represent?

Satellites can measure sea surface temperature with an uncertainty of 0.3 ºC and surface wind with an uncertainty of 1 m/s.
- Remote Sensing of European Seas - V.Barale, M.Gad

~10^{-1}

To represent data with this level of uncertainty, using **half-precision (16-bit) should be enough**. Instead, double precision (64-bit) is the standard.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Universitat Autònoma de Barcelona

Reduced Precision Emulator

- Fortran … ry Physics group, v…

Overview

The library contains a derived type: `rpe_var`. This type can be used in place of real-valued variables to perform calculations with floating-point numbers represented with a reduced number of bits in the floating-point significand.

Basic use of the reduced-precision type

The `rpe_var` type is a simple container for a double precision floating point value. Using an `rpe_var` instance is as simple as declaring it and using it just as you would a real number:

```
TYPE(rpe_var) :: myvar

myvar = 12
myvar = myvar * 1.287   ! reduced-precision result is stored in `myvar`
```

Controlling the precision

The precision used by reduced precision types can be controlled at two different levels. Each reduced precision variable has an `sbits` attribute which controls the number of explicit bits in its significand. This can be set independently for different variables, and comes into effect after it is explicitly set.

```
TYPE(rpe_var) :: myvar1
TYPE(rpe_var) :: myvar2

! Use 16 explicit bits in the significand of myvar1, but only 12 in the
! significand of myvar2.
myvar1%sbits = 16
myvar2%sbits = 12
```

Implementing the emulator

- After some months lost manually implementing the emulator…
  a Python tool to automate the process was created.

Oriol Tintó-Prims / implement-rpe-to-NEMO

README.md

RPE implementation tool

This tool is intended to automatize the processes of implementing the Reduced Precision Emulator to a computational model.

Even it was first designed to be used for the **NEMO** model, it has evolved to be useful for **any computational model written in FORTRAN**.

The script automatically:

- Replace REAL variables.
- Track dependencies between variables.
- Fix parameter declarations.
- Fix routine/function calls.
- Fix WRITE and READ statements.
- Add module to read precisions from namelist.
- Add precision assignment for each variable.

It can be also used to:

- Generate namelists for different cases.

RPE in NEMO:
What we can do with it?

With a **single binary**, we can specify the number of significant bits used for each real variable declaration within the code through a **namelist**.

```
! namelist variable precisions
&precisions
emulator_variable_precisions(1) = 10    ! Variable:        ad_u Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(2) = 10    ! Variable:        ad_v Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(3) = 10    ! Variable:     ad_zeta Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(5) = 10    ! Variable:      pc_r2d Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(6) = 10    ! Variable:       r_r2d Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(7) = 10    ! Variable:      br_r2d Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(8) = 10    ! Variable:       p_r2d Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(9) = 10    ! Variable:      bp_r2d Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(30) = 10   ! Variable:        dTdz Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(31) = 10   ! Variable:        dSdz Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(32) = 10   ! Variable:    ad_gradP Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(33) = 10   ! Variable:      ad_phi Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(34) = 10   ! Variable:   ad_gradPx Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(35) = 10   ! Variable:   ad_gradPy Routine:     ad_balance_tile  Module:    ad_balance
emulator_variable_precisions(36) = 10   ! Variable:        ad_A Routine:      ad_bc_r2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(37) = 10   ! Variable:        ad_A Routine:      ad_bc_u2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(38) = 10   ! Variable:        ad_A Routine:      ad_bc_v2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(39) = 10   ! Variable:        ad_A Routine:    ad_dabc_r2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(40) = 10   ! Variable:        ad_A Routine:    ad_dabc_u2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(41) = 10   ! Variable:        ad_A Routine:    ad_dabc_v2d_tile  Module:      ad_bc_2d
emulator_variable_precisions(42) = 10   ! Variable:        ad_A Routine:  ad_bc_r2d_bry_tile  Module:   ad_bc_bry2d
emulator_variable_precisions(43) = 10   ! Variable:        ad_A Routine:  ad_bc_u2d_bry_tile  Module:   ad_bc_bry2d
emulator_variable_precisions(44) = 10   ! Variable:        ad_A Routine:  ad_bc_v2d_bry_tile  Module:   ad_bc_bry2d
emulator_variable_precisions(45) = 10   ! Variable:        ad_A Routine: ad_conv_r2d_bry_tile Module: ad_conv_bry2d
emulator_variable_precisions(46) = 10   ! Variable:     ad_Awrk Routine: ad_conv_r2d_bry_tile Module: ad_conv_bry2d
emulator_variable_precisions(47) = 10   ! Variable:       ad_FE Routine: ad_conv_r2d_bry_tile Module: ad_conv_bry2d
emulator_variable_precisions(48) = 10   ! Variable:       ad_FX Routine: ad_conv_r2d_bry_tile Module: ad_conv_bry2d
emulator_variable_precisions(49) = 10   ! Variable:        Hfac Routine: ad_conv_r2d_bry_tile Module: ad_conv_bry2d
```

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

RPE in NEMO: What we can do with it?

With a **single binary**, we can specify the number of significant bits used for each real variable declaration within the code through a **namelist**.

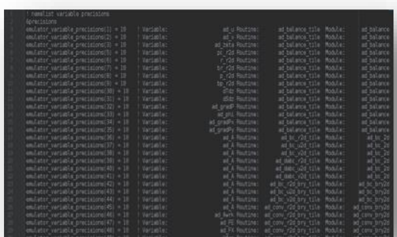

Precision namelist

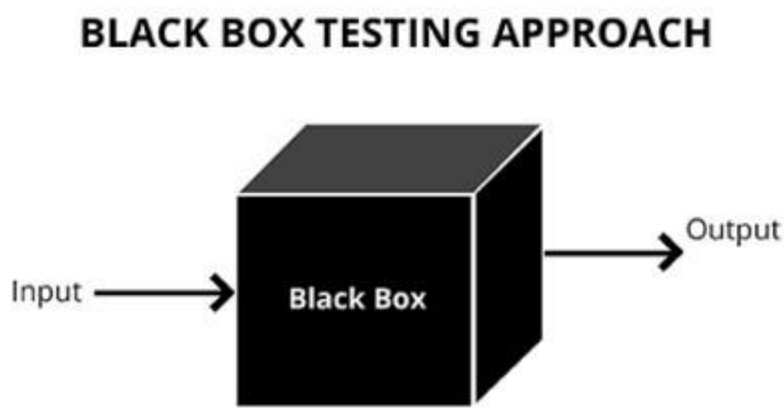

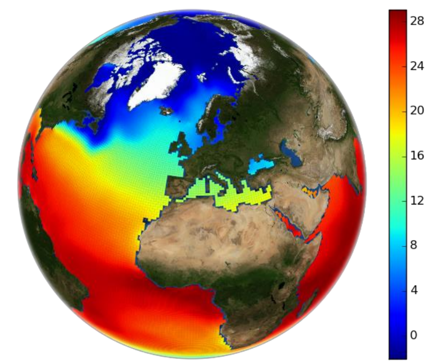

Model outputs

Verifying a non-linear model a simple example

- A simple example:
  - Lorenz system

$$
\begin{aligned}
\frac{\mathrm{d}x}{\mathrm{d}t} &= \sigma(y - x), \\
\frac{\mathrm{d}y}{\mathrm{d}t} &= x(\rho - z) - y, \\
\frac{\mathrm{d}z}{\mathrm{d}t} &= xy - \beta z.
\end{aligned}
$$

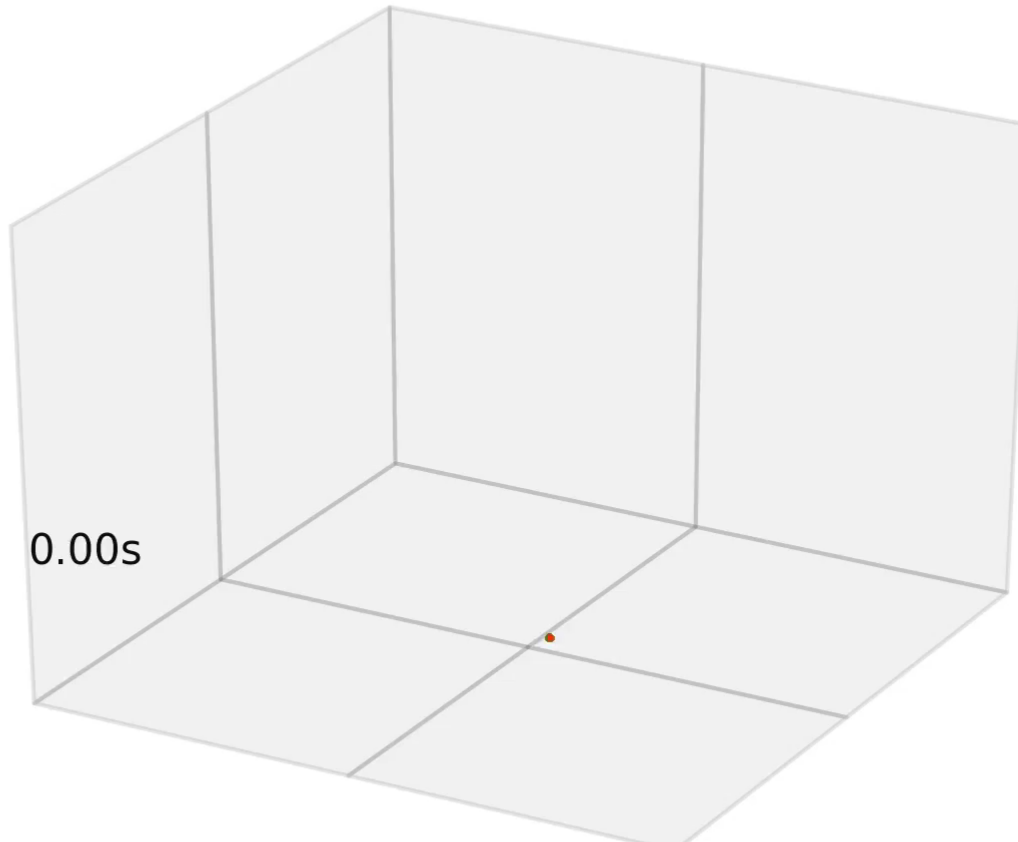
0.00s

Verifying a non-linear model
a simple example

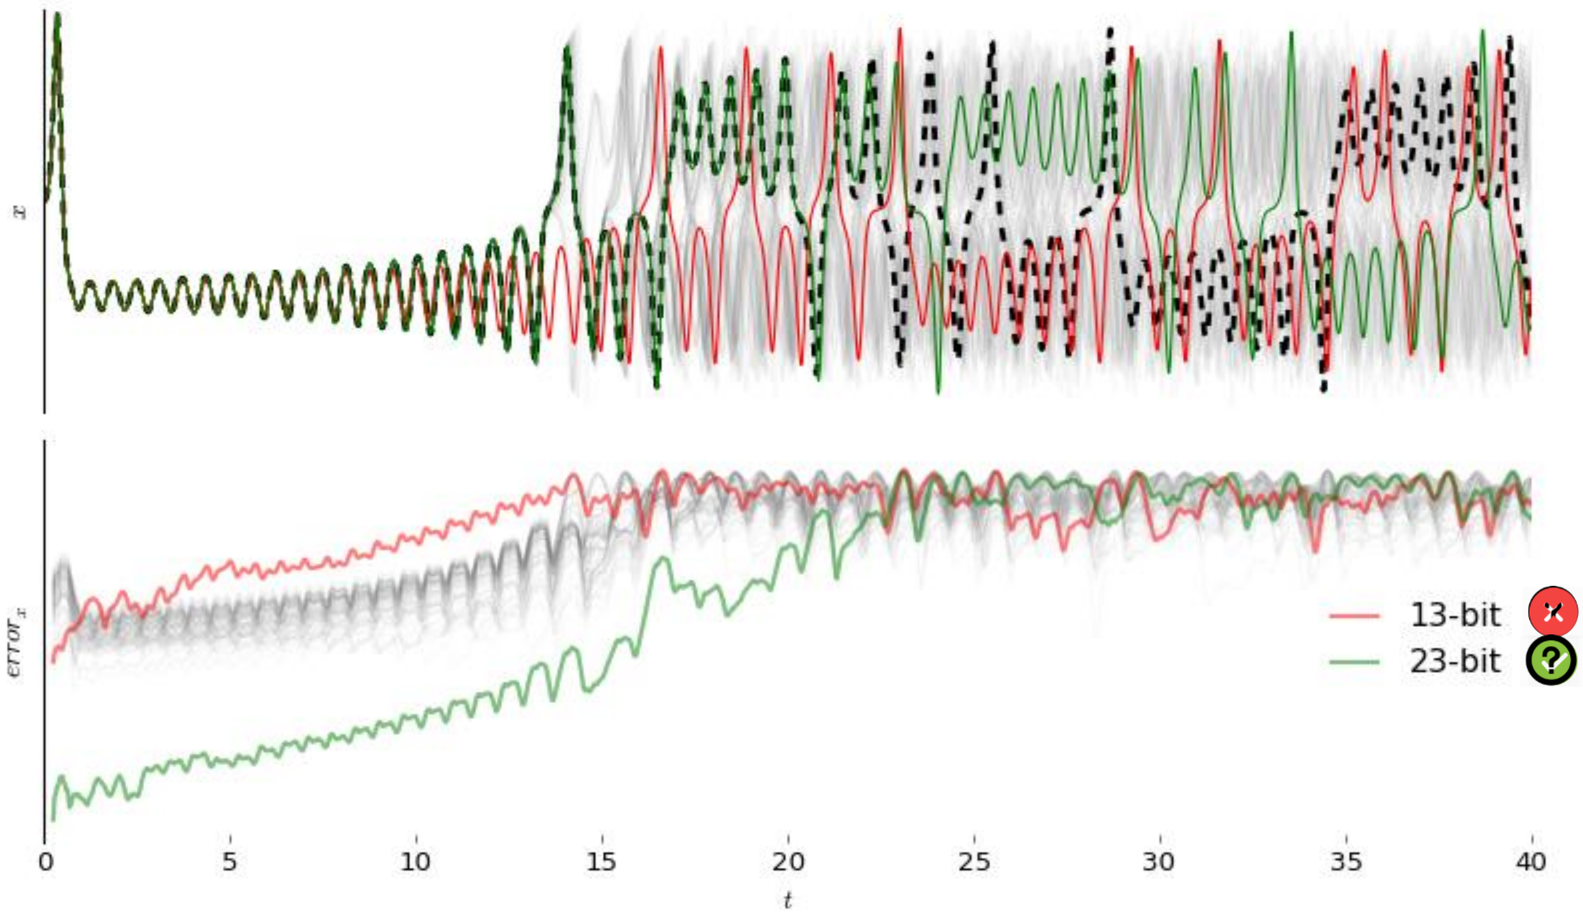

Verifying NEMO

- Initial conditions perturbed with white noise in the 3D temperature field.
- Evaluating 53 output variables.

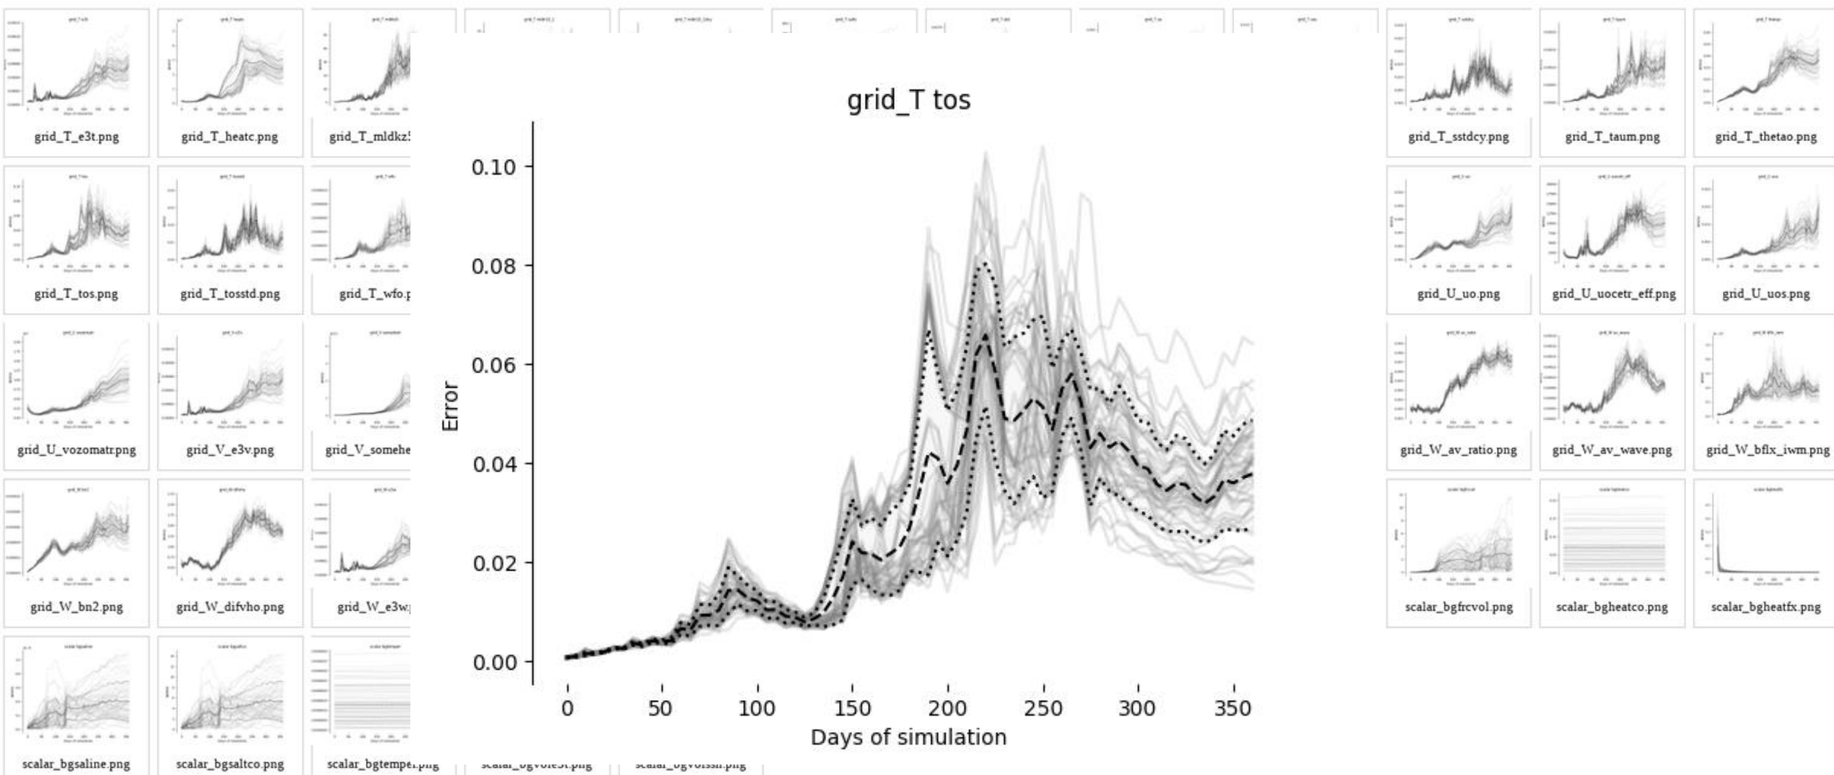

Verifying NEMO

- Example: Compiling with **–xCORE-avx512**

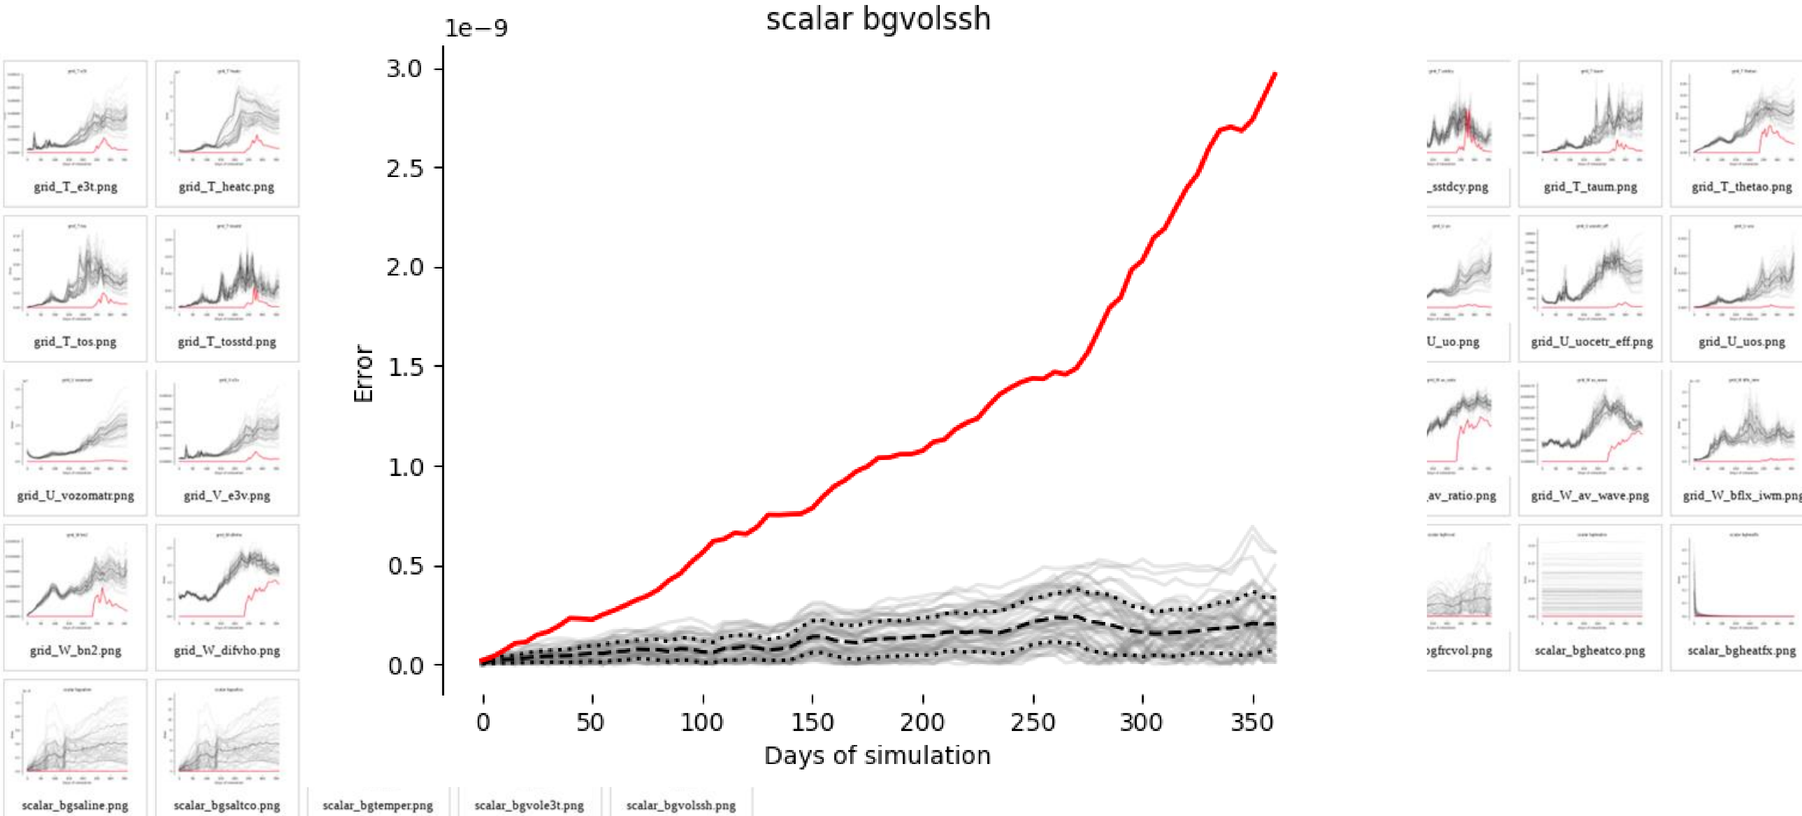

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4] ?   [5 6 7 8 9] ?

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4] [5 6 7 8 9]

[0 1 2] [3 4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4]

[5 6 7 8 9]

[0 1 2]

[3 4]

[3] [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4] [5 6 7 8 9]

[0 1 2] [3 4]

[3] [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4]   [5 6 7 8 9]

[0 1 2]   [3 4]

[3]   [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4]   [5 6 7 8 9]

[0 1 2]   [3 4]

[3]   [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4]   [5 6 7 8 9]

[0 1 2]   [3 4]

[3]   [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4] [5 6 7 8 9]

[0 1 2] [3 4]

[3] [4]

Analysis Algorithm

[0 1 2 3 4 5 6 7 8 9] ✓

[0 1 2 3 4] ✓ [5 6 7 8 9] ✓

[0 1 2] ✓ [3 4] ✓

[3] ✓ [4] ✗

↓

Variable 4 must be kept in double-precision.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

More info: How to use mixed precision in ocean models. https://www.geosci-model-dev-discuss.net/gmd-2019-20/

Analysis algorithm:
How is it implemented?

Implementation done in Python:

- class Job(variable_set):
  - Submit job to remote machine.
  - Check job status.
  - Evaluate success.
  - Expand subgroups.
  - Check subgroups.

[0 1 2 3 4 5 6 7 8 9]

[0 1 2 3 4] [5 6 7 8 9]

PENDING

RUNNING

SUSPENDED

SUCCESS

FAIL

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Precision Analysis

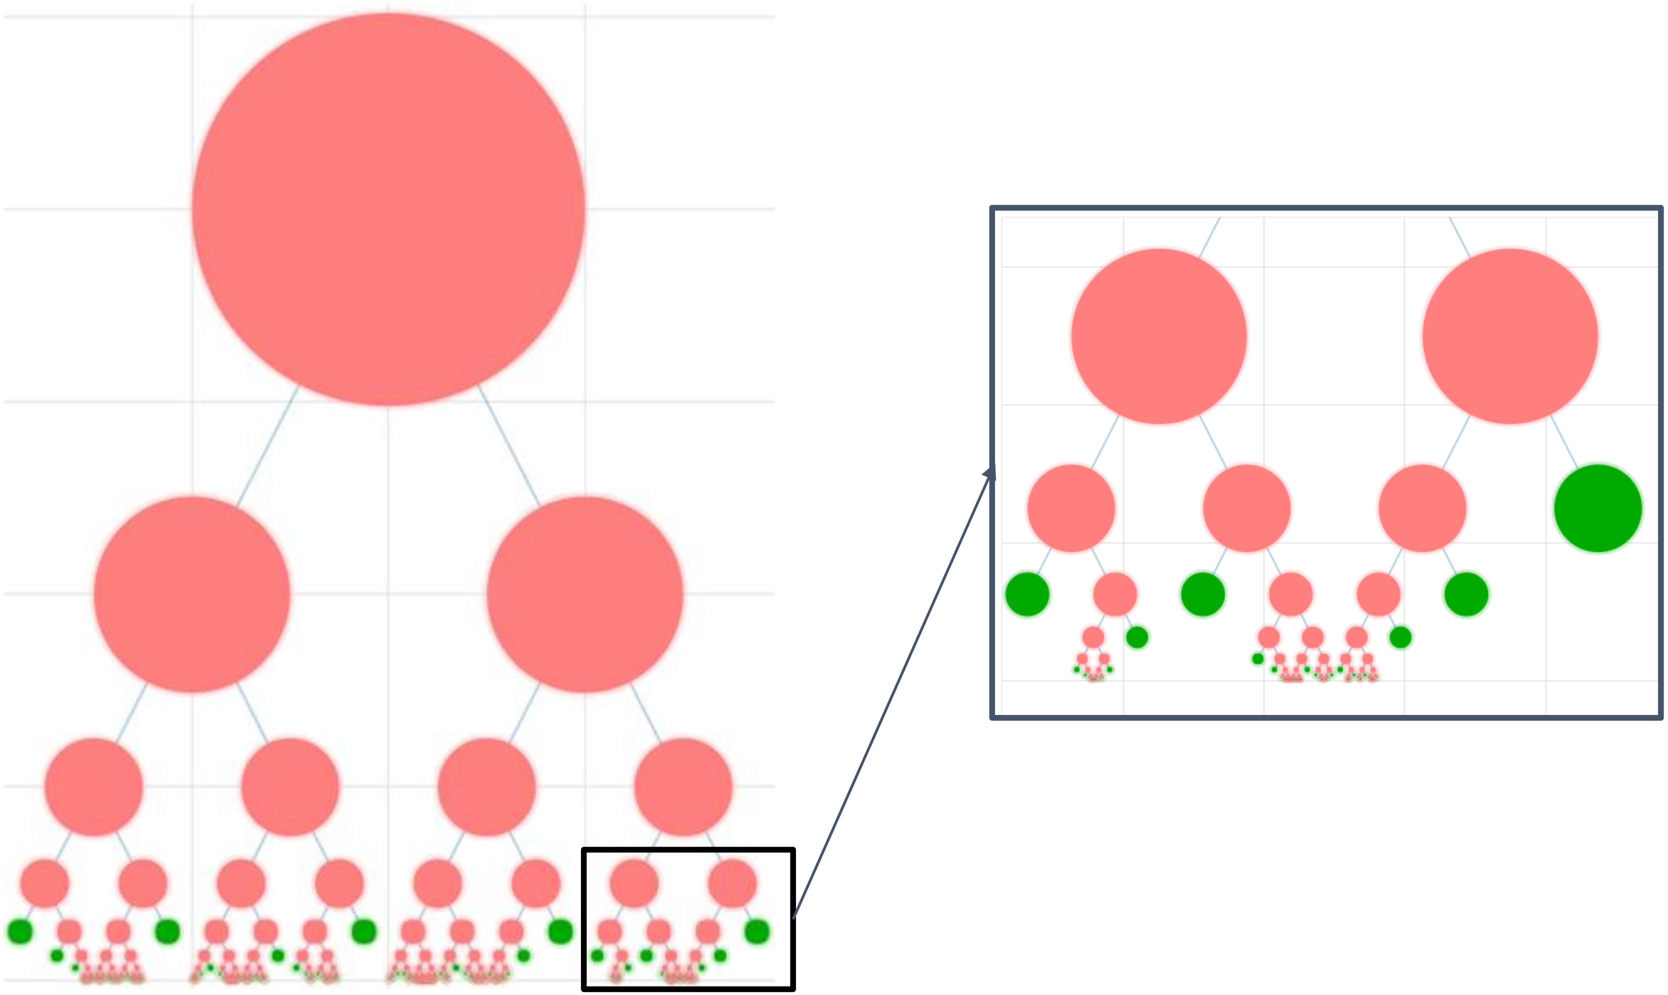

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Results

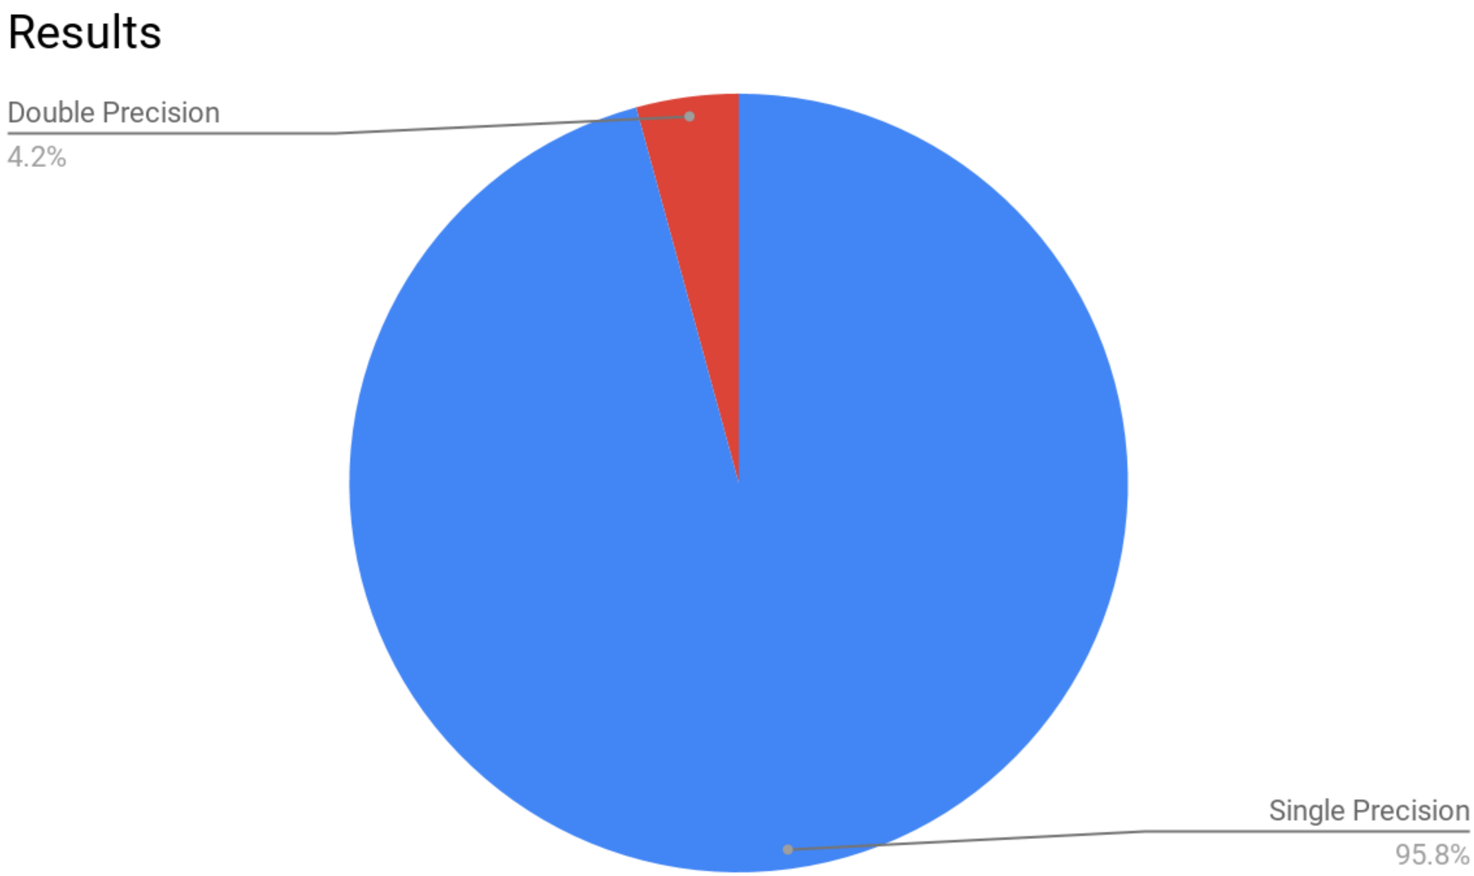

Results

- Verifying results.

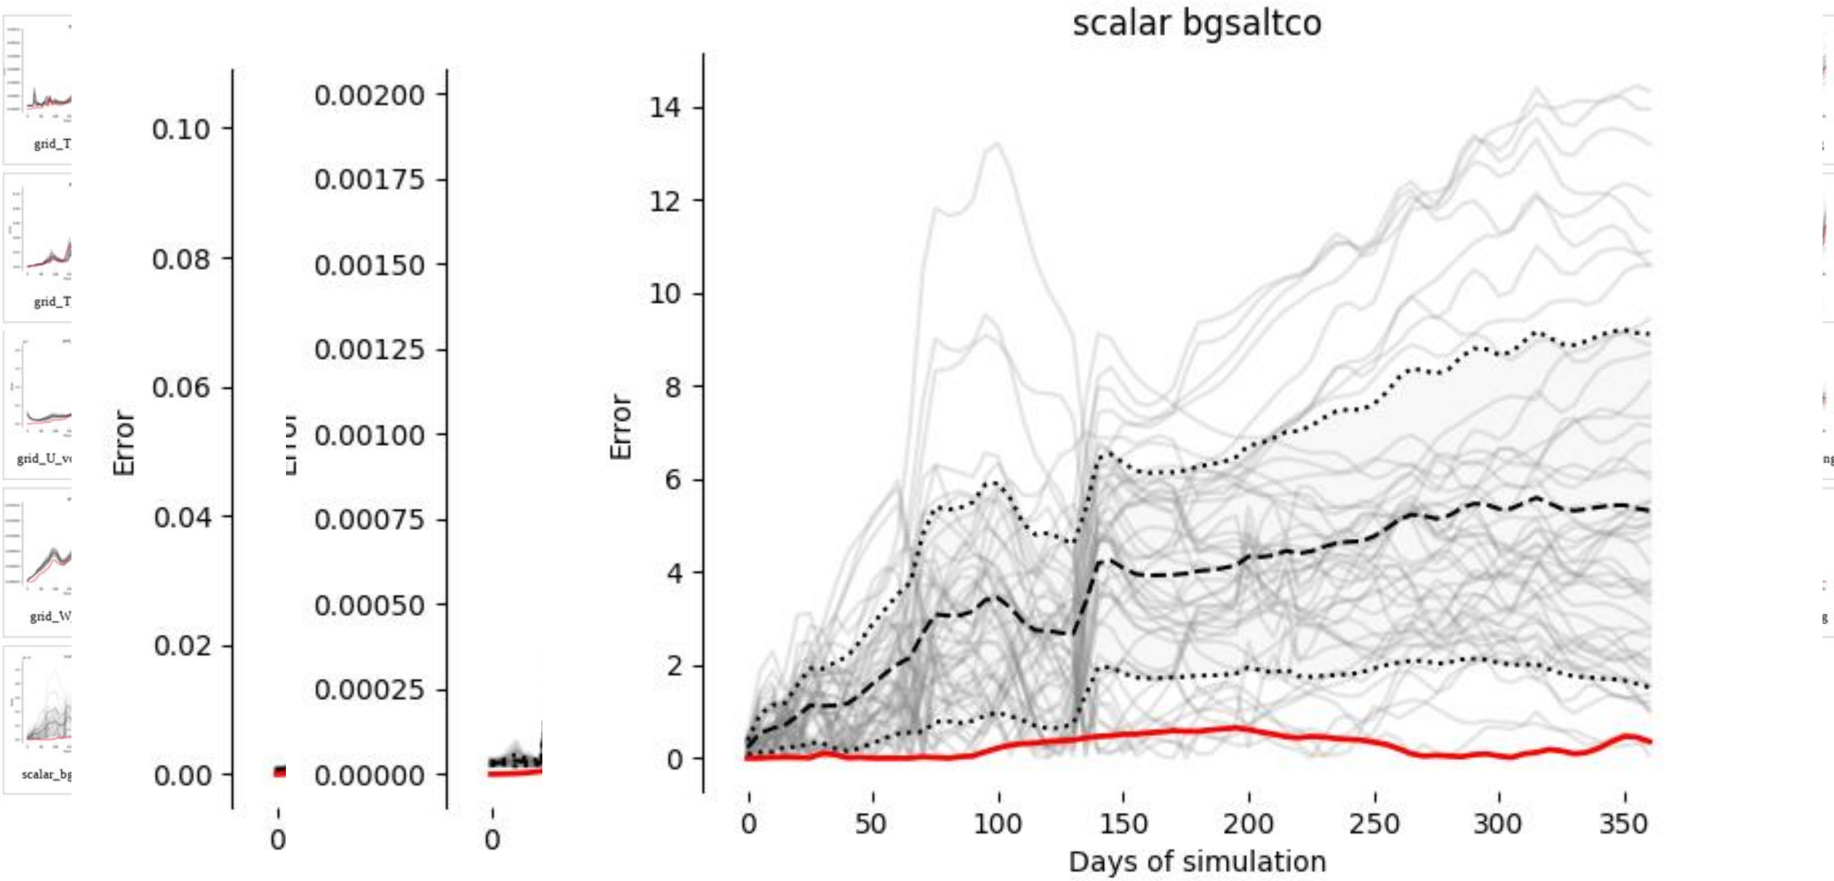

Supercomputing Center
Centro Nacional de Supercomputación

Contributions

- Performance analysis and optimization
  - Detecting, analyzing and optimizing MPI communication bottlenecks in Earth System models - Published at Journal of Computational Sciences
- Domain decomposition optimization
  - Optimizing domain decomposition in an ocean model: the case of NEMO - Presented at ICCS 2017, published at the proceedings.
- Verifying non-linear models
  - Discriminating accurate results in nonlinear models. - Submitted to WCES @ HPCS 2019
- Numerical precision analysis for Earth Science models
  - How to use mixed precision in Ocean Models - Final stages of revision at Geoscientific Model Development

Future Plans

- Finish the thesis during this academic year.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

caos

UAB
Universitat Autònoma de Barcelona

Thank you

oriol.tinto@bsc.es